


④

DTIC FILE COPY

AD-A224 689

John A. Bateman

University  
of Southern  
California



From Systemic-Functional Grammar to  
Systemic-Functional Text Generation:  
Escalating the Exchange

DTIC  
ELECTE  
JUL 31 1990  
S E D

DISTRIBUTION STATEMENT A

Approved for public release;  
Distribution Unlimited

INFORMATION  
SCIENCES  
INSTITUTE



213/822-1511

4676 Admiralty Way/Marina del Rey/California 90292-6695

## REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified			1b. RESTRICTIVE MARKINGS		
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION / AVAILABILITY OF REPORT		
2b. DECLASSIFICATION / DOWNGRADING SCHEDULE			This document is approved for public release; distribution is unlimited.		
4. PERFORMING ORGANIZATION REPORT NUMBER(S) ISI/RR-89-220			5. MONITORING ORGANIZATION REPORT NUMBER(S) -----		
6a. NAME OF PERFORMING ORGANIZATION USC/Information Sciences Institute		6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION -----		
6c. ADDRESS (City, State, and ZIP Code) 4676 Admiralty Way Marina del Rey, CA 90292-6695			7b. ADDRESS (City, State, and ZIP Code) -----		
8a. NAME OF FUNDING / SPONSORING ORGANIZATION AFOSR DARPA		8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER MDA903-87-C-0641 F49620-87-C-0005		
8c. ADDRESS (City, State, and ZIP Code) --over--			10. SOURCE OF FUNDING NUMBERS		
			PROGRAM ELEMENT NO. -----	PROJECT NO. -----	TASK NO. -----
			WORK UNIT ACCESSION NO. -----		
11. TITLE (Include Security Classification) From Systemic-Functional Grammar to Systemic-Functional Text Generation: Escalating the Exchange(Unclassified)					
12. PERSONAL AUTHOR(S) Bateman, John A.					
13a. TYPE OF REPORT Research Report		13b. TIME COVERED FROM _____ TO _____		14. DATE OF REPORT (Year, Month, Day) 1990, April	
15. PAGE COUNT 17					
16. SUPPLEMENTARY NOTATION An earlier version of this paper was presented at the 1988 American Association of Artificial Intelligence Workshop on Text Planning and Realization, held in St. Paul, Minnesota, August 1988.					
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)		
FIELD	GROUP	SUB-GROUP			
09	02		Computational systemic linguistics, text generation, text planning		
19. ABSTRACT (Continue on reverse if necessary and identify by block number)					
<p>The boundaries between planning and realization within generation are a standard division in natural language processing, although they no longer seem so secure. New territorial divisions appear necessary, but it is still unclear where the borders are best drawn. This report shows that systemic-functional linguistics (SFL) offers a rich body of linguistic work concerned precisely with issues that are coming to the forefront in natural language processing. This work provides theoretically well-motivated and thorough guidance in an area where text planning and generation are still lacking in experience: the treatment of text in context for the purposes of communication and social interaction. Many of the issues raised in text planning are also addressed within the SFL tradition; SFL can provide an extremely detailed and beneficial map of the territory through which we now need to move.</p>					
20. DISTRIBUTION / AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input checked="" type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION Unclassified		
22a. NAME OF RESPONSIBLE INDIVIDUAL Victor Brown Sheila Coyazo			22b. TELEPHONE (Include Area Code) 213/822-1511		22c. OFFICE SYMBOL

Unclassified

**SECURITY CLASSIFICATION OF THIS PAGE**

(8c continued)


Air Force Office of Scientific Research  
Bolling Air Force Base, Building 410  
Washington, DC 20332

Defense Advanced Research Projects Agency  
1400 Wilson Boulevard  
Arlington, VA 22209

Unclassified

**SECURITY CLASSIFICATION OF THIS PAGE**

University  
of Southern  
California



John A. Bateman

From Systemic-Functional Grammar to  
Systemic-Functional Text Generation:  
Escalating the Exchange

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input checked="" type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	



INFORMATION  
SCIENCES  
INSTITUTE



213/822-1511

4676 Admiralty Way/Marina del Rey/California 90292-6695

## Acknowledgments

Discussions with Christian Matthiessen and Cécile Paris significantly improved the shape, content, and rhetorical force of this paper.

## Note

Until recently it has sometimes been rather difficult to obtain reports on the results of basic research undertaken within SFL. This situation is now changing, with many collections of previously unpublished papers finding their way into print. In order to speed further this dissemination of information, interested parties are invited to contact the author at USC/Information Sciences Institute for pointers to where relevant literature may be found.

# 1 Introduction and Orientation

As guidelines for this workshop, position papers have been asked to address a small number of issues concerning text 'planning', 'realization', and the relationships and boundaries between these. Although a standard division in natural language processing, dating back to Thompson's (1977) distinction between strategies and tactics and beyond, the boundaries between planning and realization within generation no longer seem so secure. New territorial divisions appear necessary but it is still unclear where the borders are best to be drawn. My concern here will be to show that there is an existing rich body of linguistic work that has been concerned precisely with issues that are now coming to the forefront in natural language processing. This work provides theoretically well-motivated and thorough guidance in an area where computational linguistics is, still, lacking in experience: the treatment of text in context for purposes of communication and social interaction.

Programs that are to employ natural language as a medium of interaction with people need theories of how language is used, what language achieves, and what the functions of language are. Systemic-functional linguistics (SFL) provides a general theory of language as meaningful action in social context that is of unrivaled breadth and internal coherence — precisely what programs that seek to use language in context need. Therefore, I will be presenting the text planning/generation task from the perspective of SFL, re-interpreting the questions posed to guide this workshop in the light this provides. We shall see not only that all the issues raised are addressed within the SFL tradition, but also that strong proposals are made for the directions that may most fruitfully be searched for implementations. While in most cases SFL does not provide ready-built solutions to the problems of generation, it can provide an extremely detailed and beneficial map of the territory through which we now need to move.

# 2 SFL and Computational Linguistics

Over the past 20 years we have seen a slow, and rather intermittent, interaction between computational linguistics and SFL (see: Mann, 1983). Even so, there has already been a tremendous pay-off for computational linguistics: each new application and incorporation of systemic principles, embodying as it does an essential commitment to the functionality of the use of language in context, has produced state-of-the-art, and state-of-the-art advancing, results in AI. These include Winograd's (1972) SHRDLU, Davey's (1974) Proteus, and Mann and Matthiessen's (1985) Nigel. Strong systemic influences have also been present in the earlier development of McDonald's (1980) MUMBLE and Kay's (1979) Functional Unification Formalism, both highly influential in current computational linguistics.

I think it is time to make some strong claims. Without them there is a danger that the real value of current implementations of a number of aspects of systemic theory, e.g., Nigel, SLANG (Patten, 1986), and the systemic analysis grammar for English (Kasper, 1988b), is going to be missed. The principle claim is this. There is a good reason why this level of success has been achieved in computational approaches to language when insights from SFL have been applied: SFL provides a conceptual and intellectual framework that cuts the phenomena of language at the joints. It provokes the questions that need to be asked if the functionalities of language are to be understood and modeled.

In scope, SFL is already unrivaled: there is significant work from phonetics and phonology to literary stylistics, ideology, and socialization; all of these facets benefit considerably from the single overarching and encompassing framework of SFL that both places them in appropriate relation to one another and foregrounds particular directions for their development. In addition, all such work simultaneously ranges from the most theoretical to the most applied; the basic methodological premise of "renewal of connection" between theory and application that is central to SFL ensures that theory development is only undertaken as a response to extensive analyses of naturally occurring language, from spontaneous conversations to extended literary texts.

These concerns place the organization, design, and realization of texts centrally. Text is taken as the basic semantic unit in SFL — the stretch of language where all aspects of the linguistic system are made manifest. Given the overlapping areas of concern, for there not to be strong implications for computational attempts to model the processes of text creation and design would be unthinkable. There are very significant implications to be drawn, and, in the rest of this paper, I will attempt to present some of these so as to provide clarifications of the issues with which this workshop is concerned.

### 3 SFL and Text Generation

Mann *et al.* (1981) noted that there was no general set of tools and resources from which researchers could approach text planning. Although the situation has improved significantly, it is nevertheless still maintained that the concept of 'text planning' is unclear in its scope and intention (e.g., Hovy, 1988). A major contributor to this state of affairs is the lack of any theoretical basis sufficiently broad as to cover the range of tasks that need to be undertaken when designing and creating text. I am claiming here that SFL can provide such a basis. In this section, I will suggest how an escalation of the exchange of techniques and theory between SFL and text planning/generation promises to enrich our understanding of the 'text planning' problem as a whole and counter a fragmentation of research topics that is increasingly evident.

One set of techniques that is being profitably pursued at this time uses the highly developed functional discriminations that make up the Nigel grammar to uncover the semantic and pragmatic distinctions that any adequate text planner will need to support. This methodology, motivated in Matthiessen (1987) and illustrated in Bateman (1988a, 1988b), may be explained as follows.

### 3.1 Using a computational systemic-functional grammar as a tool for uncovering text planning abstractions

Until the recent resurgence of 'functionalism', mainstream linguistics has been strongly influenced by a philosophical tradition of language study that has concentrated primarily upon 'propositional content'. This tradition has also been significant in shaping a common construal of the text generation task in terms that sharply distinguish 'what' you want to say (commonly restricted to 'propositional content') from 'how' you say it. However, with the attempt to deal with more natural and sophisticated uses of language, text generation is being led away from this modularization of the task: kinds of 'meaning' in addition to the purely propositional are coming into focus. For example, computational natural language processing systems are now uncovering aspects of meaning that are concerned with the interactions between participants in the speech event and the personal attitudes and evaluations of those participants. Work here includes that on tailoring texts to their intended audience (e.g., Paris, 1987), and on 'slanting' texts to achieve interpersonal stances (e.g., Hovy, 1987; Sanford and Roach, 1987).

Standard SFL, in fact, posits three distinct types of meaning that language in use conveys: *ideational* meaning, concerned with the representation of experience; *interpersonal* meaning, concerned with the intrusion of the speaker into the social situation of speaking and with the relationship of the speaker to his/her audience; and *textual* meaning, concerned with making the language used responsive to its particular context of use. The growing awareness within computational linguistics of this latter aspect of meaning also is clear in the many approaches to text organization currently being developed (e.g., McKeown, 1982; Grosz and Sidner, 1986; Mann and Thompson, 1987), and in the very existence of this workshop.

The Nigel grammar offers resources for expressing meanings across all three of these areas of meaning, or *metafunctions*. Indeed, this is one view of precisely what a grammar is and why it is necessary: by means of the complex and tightly bound structural relationships that grammar makes possible, many distinct meanings may be woven together to form a single, 'polyphonic' structural whole that nevertheless maintains those meanings in a recoverable form (cf. Halliday, 1978). The fine functional discriminations made by the resources of the grammar also motivate an informal set of abstractions that partially specify the input that any user of these resources will need to provide in order to control them. These abstractions have already been used to



good effect in one subpart of the ideational area, the 'experiential' (Moore and Arens, 1985; Bateman, Moore, Kasper, and Whitney, 1989); current research is extending the technique to the investigation of another subpart of the ideational, the 'logical' (e.g., Bateman and Paris, 1989), and of the interpersonal (e.g., Bateman, 1988b) and textual areas (e.g., Matthiessen and Bateman, in preparation). This is bringing valuable additional constraints to bear upon how these types of meaning may be articulated computationally and offers a way of bootstrapping our knowledge up into the higher levels of semantics and pragmatics that present text planners/generators need.

Given the metafunctional organization of linguistic strata as adopted in SFL, computational research upon interpersonal and textual types of meaning can be clearly positioned and inter-related within the overall task of text creation. Consequently, important properties of these types of meaning can be motivated from SFL directly. For example, SFL theory makes interesting claims about the typical ways that meanings from the three metafunctions are realized in language (cf. Matthiessen, 1988b). The experiential subtype of ideation strongly favors 'building block', constituency-style organizations; the other subtype, 'logical', is responsible for the dependency-style organization; interpersonal meanings strongly favor 'prosodic' organizations that persist over stretches of text; and textual meanings favor 'pulse'-style organizations that may cut across the constituency and prosodic strands of organization. Text generators are beginning to feel the need to be sensitive to these distinctions. For example, while standard constituency analyses illustrate the first mode of meaning, new approaches such as Hovy's suggestion of 'in-line' planning for 'pragmatic and stylistic' goals (Hovy, 1988) offer approximations to the prosodic style of organization and, although not yet fully in focus in computational work, treatments of intonation (e.g., Pierrehumbert, 1980) and of the creation of constituents that are apparently 'discontinuous' (e.g., "That bed hasn't been slept in by anyone for years") illustrate pulse.

The difference, then, in moving into these 'less referential' styles of meaning with the background of SFL theory in place, is that we will no longer be surprised by the kinds of phenomena that we will initially encounter. Furthermore, we are better able to relate the variety of tasks and abstractions that need to be developed both to one another and to work already in place.

### **3.2 The relation of computational systemic-functional grammar to systemic-functional linguistics proper**

Although Nigel is already a valuable tool for use in text planning/generation, it in no way represents the full extent of what could be gained from adopting a systemic-functional base for language generation. SFL takes a stratal view of language, and the theory as a whole is concerned with each of these strata, their development, and

their interrelationships. Grammar is just one of these strata, and it is here that certain aspects have been implemented computationally to construct the computational systemic-functional grammar Nigel. Importantly, many aspects of the theory have *not* been implemented at this time. Nigel is a computational *approximation* to grammar as seen by systemic linguistics — and in certain respects, a far from complete one.

The principal reasons for the distinction that needs to be drawn between current computational systemic-functional grammar (SFG) and SFG proper are as follows.

1. Our understanding in terms of systemic theory is 'pre-computational': there is an understanding of what is involved, what phenomena are concerned, how the theory can be tested and developed, etc., but this has not been brought to the level of algorithmic specifiability.
2. The phenomena for which the SFL constructs are intended are not those that have been relevant, necessary, or sufficiently foregrounded in machine-based language generation and understanding.

The latter reason is fast changing; the former reason will persist as long as the effort of achieving computational specifiability is not made.

It is, then, essential to realize that the Nigel grammar is an implementation of certain aspects of SFG only. The Nigel grammar is both an *approximation* — in that many mechanisms are not addressed and its coverage is limited (relative to Halliday (1985a), for example, although not with respect to most other generation grammars in existence), and, very important, it is a *theoretical refinement* — in that the mechanisms implemented computationally are specified at a level of detail far beyond that achieved within non-computational SFG.

However, it is equally important to realize that the Nigel implementation of SFG is nevertheless still crucially shaped by the unimplemented background of systemic theory that is its foundation and source. Many of the design decisions of the grammar *only make sense* when the noncomputational theoretical context is considered. A treatment of some grammatical phenomenon in one way rather than another may not be judged on local criteria alone. The contribution and fit of the analysis to the rest of the grammar and the implications of the analysis at other strata must also be taken into account.

What is really significant about Nigel, therefore, is the beginning it offers on *expanding* the range of SFL theoretical constructs that may be implemented computationally. Although these then become available as resources for designing computational systems in general, the particular design of Nigel, as drawn from the well of SFL theory and as the principal supplier of computational implementations of SFL constructs, can *only* be maximally utilized by expansion in this way.

This is the key I am suggesting to approaching the issues of text planning and generation. Working out from the solid basis of the partial implementation of SFG that Nigel provides leads us directly to questions and issues of text planning, while still maintaining firm contact with the linguistic phenomena in terms of which texts are created. With the existence of Nigel, others areas of SFL can now be approached with views to computational implementation, more detailed specification, and subsequent incorporation as widely applicable resources for natural language processing tasks in general.

### 3.3 Text, discourse, and register

Formulating text structure theories is clearly an area where there should be a much greater co-operation and exchange of ideas between SFL and text planning/generation. Recent computational work has considerably advanced the understanding of text organization and how it may be controlled computationally (e.g., McKeown, 1982; Grosz and Sidner, 1986; Mann and Thompson, 1987). But there is also a significant body of work in SFL on precisely the issues of text organization and discourse (including work such as Hasan, 1978; Berry, 1981; Bateman, 1985; Butler, 1985; Halliday, 1985b; Martin, 1986; Lemke, 1987; Ventola, 1987; and many others).

The two approaches need to be related. For example, McKeown's (1982) text schemata, and the developments that have been made of these since, show similarities to the work of Hasan (1978) on *Generic Structure Potential* (GSP); Matthiessen (1988a) charts the correspondence in some detail. We can draw an analogy here with the relationship between SFG and computational SFG. Developments such as McKeown's provide computationally explicit models of text organization that go beyond the level of detail found in systemic treatments of text, and the theory of GSP makes certain predictions about abstract mechanisms that may be useful for further theory development.

For example, Paris (1987), in his extension of McKeown's schemata into new domains, found the need to define additional schema-types. This is predicted by the SFL notion of 'register'. According to SFL, individual possibilities for text organization are drawn from a developing classification of text types, or 'genres', which relate specific components of the use of instances of language (known collectively as registers) with particular aspects of text organization (GSP) and grammar. It is therefore possible to make predictions concerning the types of text structures that will occur based upon the functions required of the language. That there will be differences in applicable schemata when different uses of language are addressed, and (to some extent) what those differences might be, are issues that the theory insists that we study — providing both a framework and a motivation for doing so.

Both the schemata and the GSP approach can therefore benefit. As was the case with Nigel, the computational treatment provides necessary formalization and detail

for extending the linguistic account, and the linguistic background suggests likely connections with other components of the linguistic system and good directions for further research. Future work on text schemata should therefore be able to make use of GSP's foundation in SFL and the relations it posits both between textual organization and the control of the grammar and between particular texts and possible text types.

### 3.4 Relationships between strata – realization and 'metaredundancy'

The question of interaction between 'planner' and 'realizer' has also been posed. We can see this now in terms of the essential relationship that holds between distinct strata in the linguistic system. This relationship has been termed *realization*, although this maintains a sense of directionality that is not at all appropriate. Perhaps better, although very much at the edge of current systemic theorizing, is Lemke's (1984) notion of *metaredundancy*. This relates patterns of commitment at one stratum to patterns of commitment at another; for example, a semantic distinction may be realized by a set of grammatical distinctions, or a style of using language (formal vs. informal, etc.) or an ideological slant may be realized by distinctive patternings of grammatical and lexical selections through a text. Neither stratum is taken as 'determining' the choices in the other: they co-occur in the language that is used.

This begins to clarify the rather unclear notions that are held in computational linguistics concerning planning and realization. First, many of the traditional components of 'planning' are simply at a different stratum in the linguistic system to those of 'realization'. To conflate them may raise a variety of conceptual difficulties: integrated planner-and-realizers (e.g., Appelt, 1982) are likely to prove unwieldy, since they are requiring planning algorithms to be knowledgeable about inappropriate levels of syntax. Second, the nondirectionality of realization suggests why possible implementations, in terms of the conduit metaphor and strict top-down planning followed by realization, are likely to prove to have limited application. Third, the complexity of the relationships over patternings that occur in those realizations which are required to handle real texts demands that any full implementation provide very sophisticated communication channels between the operations of the related strata: either one can take the initiative, and either one can follow. We are now beginning to see attempts to provide for this kind of behavior in multi-level (e.g., combined syntax and semantics) unification-based systems and in calls for 'interleaved' planning (e.g., Hovy, 1988). These approaches are not, however, typically grounded with respect to any particular linguistic theory, and so then find the problem of interaction rather less constrained than it need be. I would claim, therefore, that design decisions concerning the kinds of interactions that are best supported between the levels and strata of a text planning system could be much more clearly motivated by an appeal to SFL.

## 4 Conclusion

In this position paper, I have attempted to suggest some of the gains that an increased interaction between SFL and text planning/generation would show for computational linguistics. Actually, I should make it clear that I see this very much as a necessarily bi-directional dialog. Although this has not been focused upon here, such an interaction is bound to be of great benefit for SFL also; this is illustrated in some detail in Kasper (1988a), Matthiessen (1988b), and Kasper, Matthiessen and Bateman (1988).

I have argued that a far greater clarity in our understanding of what is being attempted in text planning/generation, and how it is to be achieved, can be reached if the conceptual framework of SFL is applied. This serves not only to relate apparently quite distinct areas of research in language processing, showing the interactions that are necessary, but also to direct research into areas of implementation/extension that are most likely to result in the kind of advanced functionality that is now being sought for computational systems that use language. Without SFL as a backdrop theory, text planning/generation is depriving itself of perhaps the richest source of information that is available concerning the nature of texts and the functionality of language.

Finally, I would like to suggest the following for the future development of text generation as a research area.

- For text generation systems design, it is hoped that SFL literature might offer some useful suggestions concerning how language functioning in a context of use may be modeled. A familiarity with related systemic work in a particular area could then facilitate the uncovering of problems and help in the formulation of research tasks.
- When designing courses for teaching computational linguistics, educators might consider whether a prominent SFL course component could benefit students' outlook on the problems involved by providing them with a rich organizing framework for viewing the wide variety of functionality exhibited by language.

## References

- [1] Appelt, D. E. (1982) Planning Natural Language Utterances to Satisfy Multiple Goals. Ph.D. dissertation, Stanford University.
- [2] Bateman, J. A. (1985) Utterances in Context: Towards a Systemic Theory of the Intersubjective Achievement of Discourse. Ph.D. dissertation, Department of AI, University of Edinburgh.
- [3] Bateman, J. A. (1988a) "Uncovering textual meanings: A case study involving systemic-functional resources for the generation of Japanese texts." Paper presented at the 4th International Natural Language Generation Workshop, Catalina, California.
- [4] Bateman, J. A. (1988b) "Aspects of clause politeness in Japanese: An extended inquiry semantics treatment." In *The Proceedings of the 26th Annual Meeting of the Association of Computational Linguistics*, Buffalo, New York. Also available as USC/Information Sciences Institute, RS-88-211, July 1988.
- [5] Bateman, J. A., Kasper, R. T., Moore, J. D., and Whitney, R. A. (1989) "The Penman Upper Model - 1989." Penman Development Note. To appear as an ISI Research Report, USC/Information Sciences Institute.
- [6] Bateman, J. A., and Paris, Cécile (1989) "Phrasing a text in terms the user can understand." In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, Michigan, 1989. Also available USC/Information Sciences Institute, RS-89-240, September 1989.
- [7] Benson, J. D., and Greaves, W. S. (eds.) (1985) *Systemic Perspectives on Discourse: Selected Theoretical Papers from the 9th International Systemic Workshop*. Norwood, NJ: Ablex.
- [8] Berry, M. (1981) "Systemic linguistics and discourse analysis: A multi-layered approach to exchange structure." In Coulthard, M., and Montgomery, M. (eds.), *Studies in Discourse Analysis*. London: Routledge and Kegan Paul.
- [9] Butler, C. S. (1985) "Discourse systems and structures and their place within an overall systemic model." In Benson, J. D., and Greaves, W. S. (eds.), *Systemic Perspectives on Discourse: Selected Theoretical Papers from the 9th International Systemic Workshop*. Norwood, NJ: Ablex.
- [10] Davey, A. (1974) *Discourse Production*. Ph.D. dissertation, University of Edinburgh. Published by Edinburgh University Press, 1979.
- [11] Grosz, B. J., and Sidner, C. L. (1986) "Attentions, intentions, and the structure of discourse." *Computational Linguistics*, 12, 175-204.
- [12] Halliday, M. A. K. (1978) *Language as Social Semiotic*. London: Edward Arnold.
- [13] Halliday, M. A. K. (1985a) *An Introduction to Functional Grammar*. London: Edward Arnold.
- [14] Halliday, M. A. K. (1985b) *Spoken and Written Language*. Deakin University Press (Language and Learning series).
- [15] Hasan, R. (1978) "Text in the systemic-functional model." In Dressler, W. (ed.), *Current Trends in Text Linguistics*. Berlin: de Gruyter.

- [16] Hovy, E. H. (1987) *Generating Natural Language Under Pragmatic Constraints*. Ph.D. dissertation, Yale University.
- [17] Hovy, E. H. (1988) "Two types of planning in language generation." In *The Proceedings of the 26th Annual Meeting of the Association of Computational Linguistics*, Buffalo, New York, pp. 179-186. Also available as USC/Information Sciences Institute, RS-88-209, April 1988.
- [18] Kasper, R. (1988a) "Systemic grammar and functional unification grammar." In Benson, J., and Greaves, W. (eds.), *Systemic Functional Approaches to Discourse*. Norwood, NJ: Ablex. Also available in USC/Information Sciences Institute, RS-87-179, April 1987.
- [19] Kasper, R. (1988b) "An experimental parser for systemic grammars." In *Proceedings of the 12th International Conference on Computational Linguistics*, Budapest, Hungary, August 22-27, 1988. Also available as USC/Information Sciences Institute, RS-88-212, June 1988.
- [20] Kasper, R., Matthiessen, C. M. I. M., and Bateman, J. A. (1988) "Systemic linguistics and natural language processing: Case studies in the exchange." Workshop held at the 15th International Systemic Congress, Michigan State University, August 8-12, 1988.
- [21] Kay, M. (1979) "Functional grammar." In *The Proceedings of the Fifth Annual Meeting of the Berkeley Linguistics Society*, Berkeley, California, February 17-19, 1979.
- [22] Kempen, G. (ed.) (1987) *Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology, and Linguistics*. Hillsdale: Lawrence Erlbaum Associates.
- [23] Lemke, J. L. (1984) *Semiotics and Education*. Toronto Semiotic Circle: Monographs, Working Papers, and Prepublications. Toronto: Victoria University.
- [24] Lemke, J. L. (1987) *The Topology of Genre: Text Structures and Text Types*. University of Sydney.
- [25] Mann, W. C. (1983) "Systemic encounters with computation." *Network*, 5, 27-32.
- [26] Mann, W. C. (1985) "An introduction to the Nigel text generation grammar." In Benson, J. D., and Greaves, W. S. (eds.), *Systemic Perspectives on Discourse: Selected Theoretical Papers from the 9th International Systemic Workshop*. Norwood, NJ: Ablex, pp. 84-95.
- [27] Mann, W. C., Bates, M., Grosz, B. J., McDonald, D. D., McKeown, K. R., and Swartout, W. R. (1981) *Text Generation: The State of the Art and the Literature*. USC/Information Sciences Institute: RR-81-101, and University of Pennsylvania: MS-CIS-81-9, December 1981.
- [28] Mann, W. C., and Matthiessen, C. M. I. M. (1985) "A demonstration of the Nigel text generation computer program." In Benson, J. D., and Greaves, W. S. (eds.), *Systemic Perspectives on Discourse: Selected Theoretical Papers from the 9th International Systemic Workshop*. Norwood, NJ: Ablex, pp. 50-83.
- [29] Mann, W. C., and Thompson, S. A. (1987) "Rhetorical Structure Theory: Description and construction of text structures." In Kempen, G. (ed.), *Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology, and Linguistics*. Hillsdale: Lawrence Erlbaum Associates, pp. 85-96. Also available as USC/Information Sciences Institute, RS-86-174, October 1986.
- [30] Martin, J. R. (1986) *English Text: System and Structure*. University of Sydney.

- [31] Matthiessen, C. M. I. M. (1987) "Notes on the organization of the environment of a text generation grammar." In Kempen, G. (ed.), *Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology, and Linguistics*. Hillsdale: Lawrence Erlbaum Associates. Also available as USC/Information Sciences Institute, RS-87-177, April 1987.
- [32] Matthiessen, C. M. I. M. (1988a) Organizing Text: Rhetorical Schemas and GSP. Working paper, USC/Information Sciences Institute.
- [33] Matthiessen, C. M. I. M. (1988b) "Representational issues in Systemic Functional Grammar." In Benson, J., and Greaves, W. (eds.), *Systemic Functional Approaches to Discourse*. Norwood, NJ: Ablex. Also available in USC/Information Sciences Institute, ISI-RS-87-179, April 1987.
- [34] Matthiessen, C. M. I. M., and Bateman, J. A. (in preparation) Uncovering the Text Base. Sydney University, Linguistics Department, and USC/Information Sciences Institute.
- [35] McDonald, D. D. (1980) Language Production as a Process of Decision-making Under Constraints. Ph.D. dissertation, Massachusetts Institute of Technology.
- [36] McKeown, K. R. (1982) Generating Natural Language Text in Response to Questions about Database Queries. Ph.D. dissertation, University of Pennsylvania.
- [37] Moore, J., and Arens, Y. (1985) A Hierarchy for Entities. Working draft, USC/Information Sciences Institute.
- [38] Paris, C. L. (1987) The Use of Explicit User Models in Text Generation: Tailoring to a User's Level of Expertise. Ph.D. dissertation, Columbia University.
- [39] Patten, T. (1986) Interpreting Systemic Grammar as a Computational Representation: A Problem Solving Approach to Text Generation. Ph.D. dissertation, Department of AI, University of Edinburgh.
- [40] Pierrehumbert, J. (1980) The Phonology and Phonetics of English Intonation. Harvard University.
- [41] Sanford, D. L., and Roach, J. W. (1987) "Parsing and generating the pragmatics of natural language utterances using metacommunication." In *The 9th Annual Conference of the Cognitive Science Society*, Seattle, Washington. Lawrence Erlbaum Associates, pp. 89-95.
- [42] Thompson, H. (1977) "Strategy and tactics: A model for language production." In *Proceedings of the 13th Annual Meeting of the Chicago Linguistics Society*, pp. 651-669.
- [43] Ventola, E. (1987) *The Structure of Social Interaction: A Systemic Approach to the Semiotics of Service Encounters*. London: Frances Pinter.
- [44] Winograd, T. (1972) *Understanding Natural Language*. Edinburgh University Press.